

Disease Outbreak Radar: A Tool for Public Health Users

Cloris Feng, Tae Yoon Lee, Derek Tam




Fig. 1. Clipping of the final Outbreak Radar dashboard tool. The primary navigation panel is composed of the heatmap on the left, the bubble map in the top center, and the tabular summary in the bottom center. Time-series case data for the selected intersection of a disease and USA division can be visualized as a line-over-bar chart under the right column, which is stacked as more pairings are accumulated.

Abstract—Provincial public health users regularly assess and determine the risk of disease outbreak for various regions within BC based on current and historical case data. This information is used to take appropriate prophylactic measures and mount the necessary response. To this end, Outbreak Radar is a dashboard designed to help these users determine the outbreak risk of diseases within these monitored regions in a clear and convenient manner. The dashboard is designed with guidance from a meta-analysis conducted in this report on public COVID dashboards that collectively present a similar but adjacent type of data and existing design recommendations for dashboards. Based on these recommendations, the iterative re-design of the dashboard prototypes are detailed that ultimately lead to the final implementation, which will be tested by BCCDC public health users in the near future and potentially integrated into their process.

Index Terms—disease outbreak, detection, public health, interaction

1 INTRODUCTION

Disease monitoring is an important function of the public health system globally. With the ongoing COVID-19 pandemic, public interest in disease monitoring has risen dramatically in the past months, and various dashboards visualizing disease count data have been made available by government agencies, public health organizations, universities, and hobbyists alike. The dashboards have been developed for multiple purposes: didactic purposes for general public and specific purposes to help public health users effectively and efficiently detect trends and make informed decisions.

In particular, a critical task for public health users is to detect disease outbreaks, where a disease outbreak is defined as “the occurrence of disease cases in excess normal expectancy” [1] COVID-19 is one of more than 60 notifiable diseases or disease categories being monitored in British Columbia, and these diseases exhibit a variety of trends, such as seasonality in meningococcal disease and a gradual decline in acute hepatitis A. Consequently, it creates a great burden on public health users, who need to consider these complex disease characteristics and trends when deciding whether a reported disease count is higher than expected. Inevitably, the detection process is both fallible and costly. To alleviate the burden on public health users, an automated method has been developed in collaboration with the British Columbia Centre for Disease Control (BCCDC) [2].

The goal of this project is to develop a visualization dashboard specifically designed for public health users to help them efficiently process the output of the automated method for the main task of

detecting disease outbreaks. The rest of the report is organized as follows. In Section 2, we will review two dashboards for detecting disease outbreaks. In Section 3, we will describe the data and task abstraction for this project. In a subsequent section, we will analyze some existing COVID-19 dashboards to identify useful aspects and idioms for designing our dashboard. In addition, we will review guidelines and recommendations that have been published by professionals and academics for both design and ethics considerations. In Section 5, we will integrate the recommended guidelines and idioms to design and implement a dashboard. We will present our results in Section 6 and discuss limitations and future work in Section 7. Distribution of work across the authors is summarised in Appendix.

2 RELATED WORK

There is a dashboard developed at the BCCDC to help the public health users detect disease outbreaks using its internally developed algorithm, which is called the Public Health Intelligence for Disease Outbreaks (PHIDO) dashboard [3]. The PHIDO dashboard only provides users with a 2D line plot to show the observed disease count cases and the predicted disease outbreak levels over a selected period of time. Similarly, in our project, we also planned to implement a dashboard to visualize disease outbreaks using the automated method mentioned in the introduction. However, in this project, we considered more aspects of the visualization of disease outbreaks,

such as visualization on disease outbreak levels with geographical context. Even sharing similar project goals, our dashboard design would have more implementation features than those in the PHIDO dashboard, which could provide public health users with visualization on not only the trend of disease case numbers and outbreak levels but also geographical information of disease outbreaks.

Besides the PHIDO dashboard, we found the Somalia Polio Room Dashboard, which was designed for disease outbreak detection as well, introduced in a case study article conducted by Kamadjeu et al. [4]. The Somalia Polio Room Dashboard aims to address the needs of timely information on cases and performance indicators of the polio outbreak and to provide decision makers with a graphical display of the needed information. After introducing the dashboard design, the authors conducted an analysis on the dashboard design and indicated that the Somalia Polio Room Dashboard is able to provide users with efficient information display and integration. Also, the authors learned some lessons from their analysis on the key factors for the dashboard design for the disease outbreaks. Similarly, in our project, we conducted an analysis on some existing COVID dashboards and design guidelines and learned lessons from the analysis, which were used in helping us come up with our dashboard designs.

3 DATA AND TASKS

3.1 Data

Disease count data consist of spatio-temporal aggregations, providing the number of disease cases over a region during successive, regular time intervals. The BCCDC disease count database is composed of weekly counts of about 60 communicable diseases in 16 regions of British Columbia. As the BCCDC database is not available to the public, we used a publicly available disease count database of the United States (US) from Project Tycho as a proxy [5].

Overall, the Tycho database was an adequate but not ideal proxy for the BCCDC database. After preprocessing the database, a clean subset was obtained with no missing values. There were 9 regions (US census/statistical divisions) with 6 diseases, each of which being collected on a weekly basis from 2009-01-04 to 2014-07-27. This translated to 290 data points for each disease in each state; this value is sufficient for the automated method (it requires at least three years of the weekly counts for reliable results). However, the number of diseases and the number of regions were much smaller than those in the BCCDC database. While the number of regions was somewhat comparable (9 vs. 16), the number of diseases was only a fraction of the desired target (6 vs. 60). In the dashboard design, we took into consideration of this key difference in cardinality and chose idioms appropriately.

We collected an additional attribute on the number of population in each region using the US Census Bureau data [6]. This attribute was needed to derive an attribute on the rate of cases, calculated as the number of cases divided by the number of population. Description of the finalized database is provided in Table 1.

Table 1. Description of the Tycho Disease Count Database

Attribute name	Attribute type	levels/range	Description
Disease	Categorical	6	Disease name

Week date	Sequential/ ordered	2009-01-04 - 2014-07-27	Start and end dates of a week in which the number of cases is collected (290 observations)
Region	Categorical	9	US census/statistical divisions
Number of cases	Quantitative	0 - 3140 (integers)	Number of disease cases collected in the week date
Number of population	Quantitative	14,469,650 - 62,382,273	Number of people the year of the week date
Rate of cases	Quantitative	0.16 - 1094.90 (per 10,000,000)	Rate of cases, calculated as the number of cases divided by number of population

We now describe the derived input and output attributes for the automated method. The automated method took three inputs: the number of cases, the week number, and two trigonometric bases. The week number is the week date starting from 2009-01-01 (so that it is equal to 1 on 2009-01-07), and it was used to account for both the long-term and short term time trends in the disease count. The trigonometric bases were used to incorporate the seasonal pattern of the disease count. Note that the method was run on each disease for each region.

The automethod method essentially provides a robust curve of the disease count by down-weighting the counts that are abnormally high. The direct output of the method is an ‘robust’ estimated number of cases at each of the week numbers. Based on the method, for each week number, the probability of obtaining the number of cases at least as extreme as the observed number of cases was computed; this is referred to as the (individual-level) p-value. Then an alert level for a disease count was derived from the p-value. The levels need to be defined by the user. For instance, the BCCDC uses four levels: none ($p\text{-value} > 0.05$), low ($0.01 < p\text{-value} \leq 0.05$), medium ($0.001 < p\text{-value} \leq 0.01$), and high ($p\text{-value} < 0.001$).

For a measure of a disease outbreak, we computed the outbreak p-value based on the last m numbers of cases, where m is specified by the user. For instance, the BCCDC uses $m=3$. Then similarly, an outbreak level was derived from the outbreak p-value according to the categorization defined by the user. Table 2 summarizes the derived attributes.

Table 2. List of Derived Attributes for and from the Automated Method

Attribute name	Attribute type	Levels/range	Description
----------------	----------------	--------------	-------------

Input attributes			
Week number	Quantitative	1-290 (integers)	Week number starting the minimum of Date
Trigonometric bases	Quantitative (cyclic)	$\cos(2\pi*7*t/365.25)$, $\sin(2\pi*7*t/365.25)$	Trigonometric bases based on the week number
Output attributes			
Estimated number of cases	Quantitative	Non-negative real valued	The number of cases estimated by the automated method
Outbreak p-value	Quantitative	0-1 (real-valued)	Probability of obtaining the sum of the three most recent number of cases at least as extreme as the sum of the three most observed number of cases based on the automated method.
Outbreak level	Categorical	User-defined	Level of outbreak specified by public health users based on the outbreak p-value.
p-value	Quantitative	0-1 (real-valued)	Probability of obtaining the number of cases at least as extreme as the observed number of cases based on the automated method
Alert level	Categorical	User-defined	Level of alert specified by public health users based on the p-value. For example, an alert level is low, medium, or high if the outbreak score is 0.01-0.05, 0.001 - 0.01, or less

			than 0.001, respectively.
--	--	--	---------------------------

3.2 Tasks

There are three main tasks that public health users are interested in:

- T1.** As the total number of diseases to monitor across the regions is insurmountable ($60*16=960$ for the BCCDC), one important task for public health users is to efficiently search which disease is at high risk of an outbreak in each region. This task is performed each time a new number of cases is reported (for instance, it is weekly for the BCCDC), contributing to a significant burden on the public health users.
- T2.** If a disease with a high outbreak level (or equivalently low outbreak p-value) is found, then they need to analyze the trend in the number cases of the disease as well as the output of the automated method, such as the estimated number of cases and corresponding alert level at each time point, so as to confirm the validity of the outbreak level. For example, the BCCDC is most concerned with the last three weekly number of cases, from which the outbreak level is computed.
- T3.** Moreover, the users need to compare the outbreak levels of a disease across the regions and investigate whether there is any spatial pattern.

4 ANALYSIS OF COVID DASHBOARDS AND DESIGN RECOMMENDATIONS

In the analysis section, we first conducted an analysis on some existing COVID dashboards, as illustrated in section 4.1. Then, we analyzed some dashboard design guidelines and recommendations in section 4.2. Finally, we summarized all the key points made in sections 4.1 and 4.2 and presented them in section 4.3.

4.1 Analysis of COVID Dashboards

As the COVID-19 pandemic develops around the world, dashboards for tracking information on COVID cases and case development trends are available online for public use. In this analysis, we reviewed a few COVID dashboards and extracted features and idioms from them, which were useful and inspiring in helping us come up with ideas on our dashboard design. Even though there were differences in tasks between COVID dashboards and ours, we still learned some useful lessons from the existing COVID dashboard implementations and planned to apply them to our design. More specifically, the lessons that we learned from this analysis will be useful in helping us address T2 and T3 in the task abstraction.

4.1.1 Overview Design on Reported Case Counts of A Specific Disease

In some existing COVID-19 dashboards [7]–[10], aggregated confirmed COVID-19 cases are presented in the form of a choropleth, such as the one shown in Fig. 2, where the numbers of reported cases across the states are scaled in the different degrees of color change. The redder a specific region on the map is, the more confirmed cases are reported in that region. This form of representation is beneficial for users to identify the regions with most reported cases by judging the hues shown on the map. However, the data in Fig. 2 is collected at the county level, which would be visually heavy for users to spot the states with most accumulated cases. In this project, we could use a choropleth to show the

aggregated cases for a specific disease across all states so that users could quickly identify the regions with most reported cases, which is relevant to the need of T3 of this project. Also, since both the proxy dataset and the BCCDC dataset were collected at the provincial level, which could lessen visual burden for users while showing the overview of a specific disease across all regions, a choropleth for reported confirmed cases could be used to provide users with information on confirmed cases of a specific disease across all regions. Since the choropleth shows case numbers of only one disease in all collected regions at a time, the differences in disease numbers and region numbers between the proxy dataset and the BCCDC dataset will not affect the display scale of the choropleth. However, one thing worth noting here is that the severity of how COVID spreads over the population within a specific region isn't addressed with the total confirmed cases. Therefore, some COVID dashboards [7], [10] provide users with choropleths of total confirmed cases by population as well, such as the one shown in Fig. 3. We could see from Fig. 2 and Fig. 3 that there is a shift of regions highlighted in red, which validates the idea that only showing a choropleth of total confirmed cases is not enough to illustrate the severity of the spread of a disease on the population in an area. Fig. 3 provides users with information on regions where the population is affected most by the disease. This could be included as an option for the dashboard users in the project implementation. Similar to the traditional choropleths, the display scale of the normalized choropleths by population will not be affected by the differences in disease numbers and region numbers of the proxy dataset and the BCCDC dataset, as data of only one disease is shown on the map at a time.




Fig. 2. Aggregated confirmed COVID cases in the United States are illustrated in a choropleth at the Johns Hopkins Coronavirus Resource Center website [10].

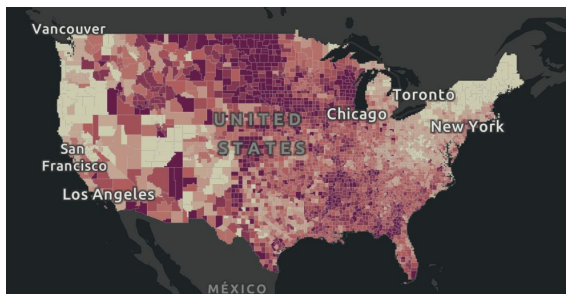


Fig. 3. Aggregated confirmed COVID cases by population in the United States are illustrated in a choropleth at the Johns Hopkins Coronavirus Resource Center website [10].

4.1.2 Overview Design on Disease Outbreak Levels of A Specific Disease

Besides choropleths, some COVID dashboards use other idioms to show total cases reported in the United States. For instance, the bubble map shown in Fig. 4. The size of bubbles on the map represents the amount of total COVID cases in a specific state. The larger the bubble is, the more reported cases are within the specific region. Also, when users hover over a specific bubble in a region, a box will appear and show the total, active, recovered, and fatal case counts in that region and each case count category is represented with a colored circle legend for better visualability. The usage of bubble size in this representation is not of high distinguishability, when the amounts of case counts do not have large differences, leading to less differences in bubble size and more difficulties in judgement on case number differences by eyes. However, in this project, a similar bubble map design could be used to show an overview of disease outbreak levels of a specific disease across all regions. In our project, we would have low/medium/high outbreak levels to classify the risk of disease outbreak in a specific region. In this case, we could use different colored bubbles, such as green/yellow/red bubbles to represent low/medium/high outbreak levels of a specific disease. Instead of having orange bubbles shown on map as illustrated in Fig. 4, we will show the outbreak level of a specific disease in a region by showing its corresponding colored bubble within the specific state. In this way, the user could quickly identify states at a high risk of disease outbreak by finding states with red bubbles shown on the map and this representation could meet the need of T3 of this project. Since the bubble map will only show the outbreak levels of a single disease across all regions on the map and there are only three defined disease outbreak levels, there will only be one of the three colored bubbles in each individual region on the map to represent the corresponding outbreak level of a region and the differences in disease numbers and region numbers between the proxy dataset and the BCCDC dataset will not affect the scalability of the idiom in both cases.



Fig. 4. An overview of total cases in the United States shown in a map at the Bing COVID-19 tracker dashboard [8].

4.1.3 Trend of Reported Case Counts

There are different idioms in COVID dashboards that show the trend of reported case counts [7]–[10]. In this section, we will analyze and extract and combine features from these idioms and find the design that is more suitable for T2. In Fig. 5, a bar chart is used to show the trend of confirmed cases reported weekly in the United States. Also, if users hover over a bar, they will have more information on the weekly change of case counts compared with the reported cases from the previous week. In contrast, Fig. 6 shows a line over bar chart, which illustrates new daily cases reported in a state with a bar chart and the 7-day moving average of the new COVID cases with a line on top. More specifically, the green segment of the line indicates a day-over-day decrease in new cases while the red segment of the

line means a day-over-day increase in new cases. If there is no change, the segment of the line will be grey.

In our project design, we could combine features from these two idioms. We could use a line over bar chart to show the trend of weekly reported case counts of a specific disease in a specific region, where the bars will represent weekly reported cases and the line will be the curve showing the trend of reported cases. Specifically, the curve will be color-coded to indicate the trend of increase, no change, or decrease of weekly reported cases similar to what is shown in Fig. 6. The green segment of the line indicates a weekly decrease in reported cases while the red segment of the line means a weekly increase in reported cases. If there is no change in case numbers, the segment of the line will be grey. When users hover over a specific bar, they will have an information box, similar to what we've seen in Fig. 5, which shows reported case counts in a specific week and the percentage of change compared to reported case amounts from the previous week. This idiom implementation is able to address the task stated in T2, while providing more details, such as percentage of change in case numbers, to users if they are interested. Since the line over bar chart will only show the trend of disease case number development of a specific disease in a specific region over a certain time period, the discrepancies in disease counts and region numbers of both the proxy dataset and the BCCDC dataset will not have an effect on the scalability of the idiom.




Fig. 5. Trend of total confirmed cases reported weekly in the US shown at the WHO website [7].



Fig. 6. Trend of new daily COVID cases reported in the state of Illinois and the trend of 7-day moving average [9].

4.2 Analysis of Design Guidelines and Recommendations

Besides conducting an analysis on COVID dashboards, we also reviewed some online articles regarding design recommendations and guidelines on how to present data in a dashboard layout. We have grouped the recommendations and guidelines into different categories shown below.

4.2.1 Presentation of Reported Case Numbers

As mentioned in [11], simply displaying reported case counts on a choropleth is not enough; instead, we should consider normalizing data by population before showing it on a choropleth or using a symbol map to show raw data instead. A choropleth by population could better show how a disease is impacting the population in a

region, which is mentioned in the analysis above as well. One thing is worth noting about the choropleth is that normal projection of a place on the map is not showing accurate regional area due to the globe shape of the Earth; instead, we should consider using the Albers equal area projection on the choropleth map, which could show a more accurate shape of a specific region on a map [12]. Moreover, a symbol map could demonstrate how severe a disease is in a region without being limited by the size of the area on the map; besides a symbol map, a dot density map is a good way to represent case numbers in a region since the density of dots in a region could reflect how severe a disease is in a region [12].

Furthermore, as mentioned in [11], [13], [14], color coding is important when showing data on a map; more specifically, we should avoid using red color as it tends to represent danger and death and could incite concerns or fears in the audience. Therefore, for continuous numbers, such as disease case counts, we should consider using a linear color scale, such as the viridis color scale, which avoids red color. For categorical data, for instance, various disease outbreak levels, we could utilize distinct colors other than red.

With the above recommendations and guidelines, we have come up with the following suggestions on our dashboard design, which caters to the need of T3. We could show normalized data by regional population on a traditional choropleth where the map is gained through Albers equal area projection. Also, we could visualize raw case counts on a symbol map or a dot density map. In this way, we could better visualize the practical impact of an outbreak of a disease on the population specific regions by looking at their corresponding densities of symbols or dots shown on the map. Since symbol maps and dot density maps will only demonstrate the severity of disease case counts of one specific disease across all regions at a time, the discrepancies in the numbers of diseases of both the proxy dataset and the BCCDC dataset will not have an effect on the scalability of the idiom. Also, all color coding used in these maps would leave out color red.

4.2.2 Scale of Representation

In [12], the author mentioned that it could cause misleading information representations when the system allows users to zoom in/out too far. Therefore, it is important to limit the zoom levels that the users are allowed to adjust to avoid data misinterpretation. In other words, we need to consider setting a limit that is logical for the granularity of the visualization. Besides zooming scales, the scale of data representation when comparing diseases needs extra caution as well, according to [14]. The author further explained that comparing two diseases is not as simple as juxtaposing two disease graphs, since there are differences in testing and treatments, outbreak levels, and raw cases of various diseases. Therefore, the graphical representations might not equally represent the actual risk of diseases, not to mention comparing the differences among these diseases. Lastly, we should also pay attention to the time scale that we set on the dashboard. According to [15], not all historical data is useful to demonstrate a viewpoint. In our dashboard, we should consider having the users to specify the time frame to be drawn and the visualization should show the relevant data within that time frame. However, in our case, we should set the rightmost data point to be the most current case count to constrain use to the intended tasks as demonstrated in the previous sections.

4.2.3 Notations on the Dashboard Design

As mentioned in [11], [16], showing testing and recovery numbers is important for rapidly changing diseases, such as the COVID. Even though it might be less relevant for our data, low base rate diseases with known treatments where patients self report to their doctors, we should keep this in mind in case we would have rapidly changing

diseases involved in our data. Also, if data is missing for a specific week, the frequency and quality of reports could vary, as mentioned in [11]. If this happens in the dataset in the future, we should flag this on the visualization with a different color to make the missing portions explicit. Moreover, we should also note the parameters and inputs of the models used thoroughly to avoid use of outputs without any context, according to [16], [17]. Lastly, as mentioned in [17], [18], dashboard designers have the responsibility to keep the dashboard accurate and up-to-date. In our dashboard design, it means that data should be easy to update on a weekly basis so that the dashboard can continue to be used.

4.3 Summary of Analyses and Design Recommendations

Here is a summary of key points mentioned in sections 4.1 and 4.2, which would address the need of T2 and T3 and some implementation design details.

- For T2:
 - Use a **line over bar chart** to show the trend of weekly reported case numbers of a specific disease in a specific state, where the bars will represent weekly reported case numbers and the line will demonstrate the trend of change of case numbers .
 - The line will be color-coded to indicate the trend of increase or decrease change of weekly case counts. The green segment of the line indicates a weekly decrease in reported cases while the orange segment of the line means a weekly increase in reported cases. If there is no change in weekly reported cases, the segment of the line will be grey.
 - When users hover over a specific bar, they will have an information box, which shows reported case counts in the week and the percentage of change compared to cases from the previous week.
- For T3:
 - Use a **symbol map** or a **dot density map** to show the weekly reported case numbers for a specific disease across all states so that users could quickly identify the regions with most reported cases.
 - Add a **choropleth** of total weekly confirmed cases by population to show the severity of the spread of a disease on the population in an area.
 - Use the Albers equal area projection for the map on the choropleth.
 - Use a **bubble map** to show the outbreak level of a specific disease in a region by showing its corresponding colored bubble within the specific region.
 - On the map, three bubbles of different colors represent low/medium/high outbreak levels of a specific disease.
 - Avoid using color red on all maps.
- Other recommendations on implementation:
 - Set a limit on the zoom level to avoid data misinterpretation.
 - Allow users to specify the time frame to be drawn.
 - Set the rightmost data point to be the most current case count.
 - Flag missing data with a different color on the dashboard.
 - Note the parameters and inputs of the model on the dashboard.
 - Allow users to update data on a weekly basis to maintain information up-to-date.

5 PROTOTYPE AND IMPLEMENTATION

In order to integrate the guidelines from our analysis in the context of this dashboard, a number of prototypes were designed and iterated upon in order to arrive at the final implementation. In this section, those prototypes will be discussed as well as the design changes that were made. Finally, details on the final dashboard will be provided in addition to a description of the expected workflow.

5.1 First Prototype

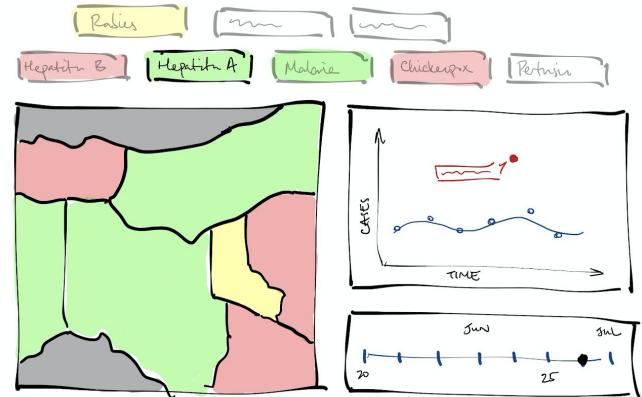


Fig. 7. Schematic drawing of first dashboard prototype.

The first prototype was made concurrently with the analysis component of this project, and was therefore created without any knowledge of health-dashboard-specific guidelines or recommendations. Driven purely by the task abstractions and the algorithm that was providing the novelty in this dashboard, the focus and key visualization naturally became the time-series visualization. This type of visualization could display the robustly fitted curve which was enabling the identification of potential outbreak timepoints. Simple visual markers and detailed call-outs were discussed in order to retain individual-point detail, but also to provide quantitative justification for outliers labeled by the model.

To manipulate the time-series, the addition of a time-window slider was also considered, which could be used to add utility through selection of a time-frame of interest. This could allow the end user to scrub through historical data and look for variable periodicity. Additionally, the timeline would provide context for the time frame being presented in the time-series as well as the complete window of data available.

To supplement the time-series, a choropleth was considered in order to provide geospatial context to the timeseries. Due to the nature of the data being explored, a spatial trend could exist where adjacent regions have similar risk levels as a result of physical transmission. While an algorithm to model this transmission is outside the scope of this project, an expert user may be able to make such an inference if the data is presented spatially.

Finally, in order to switch between diseases of interest, a tab-like structure was considered using color-coded buttons. These would serve as selection tools, but would also act as a quick visual indicator of which diseases are at high risk for outbreak.

5.2 Second Prototype

The main shortcoming that was found with the first prototype was that it was not immediately obvious which regions were at risk for which diseases. While the buttons provided a general risk summary for all regions with data, it was not possible to tell which individual regions were of concern without selecting a disease first. Furthermore, this workflow of selecting a disease, then a region, was not conducive to exploring trends between diseases for one specific region. To do so would require selecting a disease, observing the risk

level for the region of interest and remembering its risk level, selecting a different disease, and making the comparison mentally.

In retrospect, this issue originated from the inside-out design process of the first prototype, where the time-series was identified first as the basic subunit of the dashboard. This led to the addition of the choropleth and the selection buttons as supplements for context, and not as key data filtering steps. As a result, the user experience of disease and region selection was designed secondary to the model visualization.

As T1 involves identifying both regions and diseases at risk for an outbreak, this was a key issue addressed in the second prototype through the addition of a heatmap. It was determined that a heatmap most succinctly summarised the relationship between two factors in a visual manner. The heatmap cells could be colored by the predetermined risk level in a categorical manner, summarising the findings of the model for all diseases across all regions. By choosing a cell, a region and a disease are selected inherently by the intersection. The two facets of this selection can be presented separately on a choropleth as before (summarising all regions for a specific disease) as well as a visual table (summarising all disease statuses within the selected region).

The intended user experience was a highly flexible integration of these three visualizations, where selection of a specific region or disease in one panel would update the other two panels to reflect the changes. This allowed for exploration driven by curiosity and discovery as new patterns and relationships are revealed through the choropleth or the visual table.

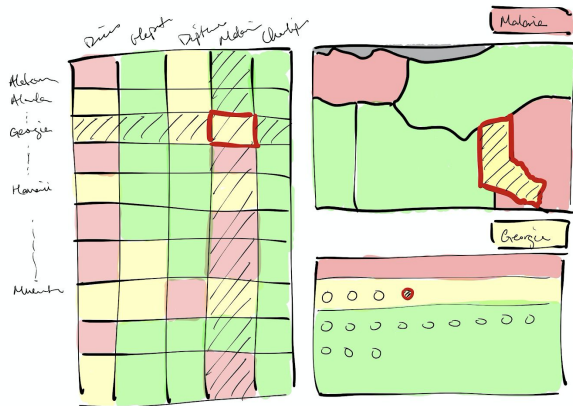


Fig. 8. Schematic drawing of second dashboard prototype's overview panels.

Once this selection has been made, it was then reasonable for the time-series to be presented as a detailed visualization of the user-specified data slice. It was imagined that a number of these time-series could be presented on a common time scale, as the user explores the various combinations and identifies interesting and important pairings for reporting or further analysis.




Fig. 9. Schematic drawing of second dashboard prototype's time-series panels. Intended to be displayed upon user selection in a stacked column.

Key questions remained after this prototype regarding the practical implementation of the choropleth, table, and the time-series. Specifically, the implementation must avoid known pitfalls when using choropleths involving user correlation of area with significance and the effect of color intensity. Presenting case data on a time-series scatter plot can also be precarious, as it can be unclear if the data is being aggregated over time, or if the cases are presented as normalized rates or as raw case numbers (both of which have their individual shortcomings).

Additionally, the decision to show the underlying predictive model's trajectory as opposed to a more generally interpretable trendline was discussed. While the former could be used to justify the significance value attributed to various outbreak risk evaluations, we believed this was outside the scope and complexity of the dashboard for its intended use. The evaluation of the model in this manner would require further manual investigation of the statistical data that will not be unavailable through this dashboard. Providing such data would be distracting from the intended use-case. In the event that the results from the model are brought into question, a more rigorous and detailed investigation should be conducted at a lower level than this dashboard, which would be encouraged by simply omitting this robust curve. In its place, a simple colored trendline to indicate gain or loss can be more powerful in showing medium- to long-term trends.

Finally, a decision was made to remove the timeline scrubbing included in the first prototype. This was fundamentally due to the intended purpose of the dashboard, which was to enable rapid assessment of outbreak risk on a weekly basis. The only relevant data-points in this calculation are the three most recent ones, and there will never be a scenario where historical data is a subject of interest. To prevent misleading the user as to the purpose and functionality of this dashboard, the timeline was determined to be superfluous.

5.3 Final Implementation

The final implementation of this dashboard was conducted in R, and was made interactive through the use of the ShinyApp package. Largely, the final implementation reflects the design of the second prototype. The heatmap is used both as a summary and selection tool, which is integrated with the geospatial and tabular visualizations. Here, the specific parameters of the idioms were changed.

In place of a classic choropleth map colored by region, the COVID dashboard analysis recommended the use of data scaled per capita presented on a bubble map. By using the size channel to

encode the case rate, and bubble color to encode the risk level, the spatial position of each bubble over its relevant region provides a clear, unbiased visual summary of the population-adjusted case rate across the covered regions.

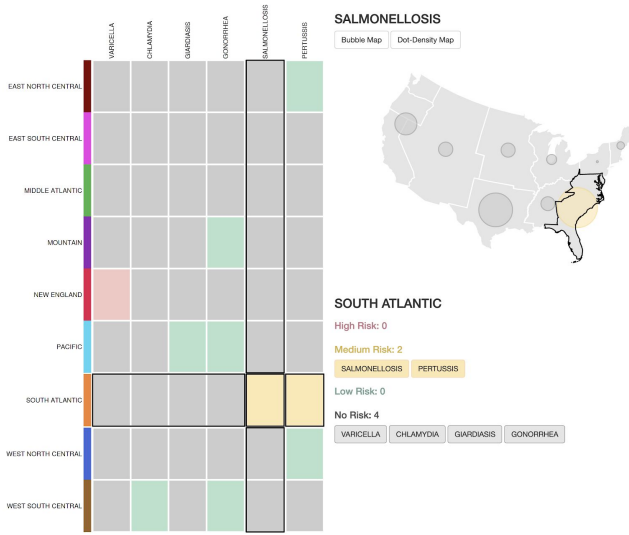


Fig. 10. Screenshot of implemented dashboard's overview panels.

Further to the recommendations, a dot density plot was added as a secondary option to supplement the bubble map. While normalized data is important for making generalized decisions that affect a region, it is equally important when working with case data to remember that each case affects a real person. To this end, many published dashboards recommend the use of a dot density plot to provide an intuition for the number of cases affecting a region. Here, color was used that is consistent with the annotations on the heatmap which may be useful in connecting the regions to the heatmap summary.




Fig. 11. Dot-density map depicting toy-data cases of salmonellosis in the USA by division. Dots do not show the precise geographic location of a case, but rather are intended to showcase density by region.

Finally, for the time-series panels, a line-over-bar chart was chosen as the idiom, combined with a colored gain/loss trendline. Based on the recommendations, this best presented the case data as “counts per time point”, showing the trendline implicitly through the change in bar height, but also explicitly using the line graph. The addition of a directional trendline color highlights runs of consistently rising or falling case numbers, which are important for intuitively understanding why a certain risk level has been assigned. Additionally, the segmented trendline utilizes angle in order to emphasize the magnitude of change between weeks, where a steep angle suggests a significant change occurred.




Fig. 12. Stacked time-series line-over-bar charts depicting historical case count and risk level for user selected disease-division pairings.

6 RESULTS

The dashboard was successfully implemented as designed, and largely feature-complete as it was conceived in the second prototype. Due to its design using R’s ShinyApp, the dashboard is portably deployable and can be updated independent of the data being visualized as it changes on a weekly basis.

Due to availability and the pressing importance of the current global pandemic, it was not possible to meet with members of the BCCDC to evaluate the performance of the dashboard. However, from peer review and internal testing, key criticisms were made regarding the dashboard’s performance and scalability. Due to the scale of the data being presented and the visual complexity of the bubble and dot-density map during rendering, the dashboard can become “sluggish” over time, especially as more time-series are requested. However, the flow for slicing data of interest was found to be intuitive and clear, as the functionality for selection and switching between views was often used without prompting or further guidance.

7 DISCUSSION

In this project, a dashboard intended to be used by public health users to determine the outbreak risk of diseases across a large region was designed and implemented. Our analyses of public COVID dashboard as well as guidelines presented by experts in these fields were used to guide the iterative design process of creating Outbreak Radar. This dashboard serves as a solid foundation inviting further improvement on the dashboard’s functionality in this capacity.

7.1 Limitations

The main limitation of this current iteration is actually a shortcoming in the selected proxy dataset from Project Tycho. Specifically, the final deployment of this dashboard will be working with 60 diseases and 16 regions, which has a large difference in cardinality. Additionally, use of color to connect the regions to the heatmap is

not conducive to such an increase in region number due to the lack of distinct color diversity in a palette of this size

7.2 Future work

Moving forward, a number of significant improvements could be made to this dashboard. Firstly, performance optimizations would be paramount as the current implementation is quite slow to render. This is exacerbated by the complexity of the maps and visualizations, but can likely be addressed by improved implementation. If this is addressed, this could likely also improve performance and scalability with a dataset that has comparable cardinality to the ultimate deployment data. This could be confirmed with the use of a better proxy that would allow for more robust testing of the dashboard and visualization design for data at scale.

Due to time constraints, not all the recommendations made in analysis were able to be fully integrated into the dashboard's design. As future work, this first iteration could be functionally improved through the addition of interactivity in the time-series bar charts. Combined with the ability to preview, reorganize, and delete these charts, the user could better curate a list of disease and region pairs that are of interest for the week. This functionality would naturally be consolidated by reporting functionality that would allow the results to be exported as a printable document.

Finally, additional exploration functionality could be added to the dashboard by allowing user-defined cutoff values for the various outbreak severity categories. While the cutoffs used here were the defaults used previously in BCCDC's PHIDO, a variable level of granularity may be desirable depending on the type of response and volume of resources available. These values would then be reflected across the dashboard and used to explore the user curated set of time-series.

ACKNOWLEDGMENTS

We would like to thank Dr. Tamara Munzner for her feedback and guidance.

REFERENCES

- [1] "WHO | Disease outbreaks," *WHO*. https://www.who.int/environmental_health_emergencies/diseases_outbreaks/en/ (accessed Oct. 22, 2020).
- [2] T. Y. (Harry) Lee, "Robust methods for generalized partial linear partial additive models with an application to detection of disease outbreaks," University of British Columbia, 2019.
- [3] BC CDC, "PHIDO (public health intelligence for disease outbreaks) for Windows: User manual," 2010.
- [4] R. Kamadjeu and C. Gathenji, "Designing and implementing an electronic dashboard for disease outbreaks response - Case study of the 2013-2014 Somalia Polio outbreak response dashboard," *Pan Afr Med J*, vol. 27, no. Suppl 3, p. 22, 2017, doi: 10.11604/pamj.supp.2017.27.3.11062.
- [5] W. G. van Panhuis *et al.*, "Contagious diseases in the United States from 1888 to the present," *N Engl J Med*, vol. 369, no. 22, pp. 2152–2158, Nov. 2013, doi: 10.1056/NEJMms1215400.
- [6] U.S. Census Bureau, "Population, Population Change, and Estimated Components of Population Change: April 1, 2010 to July 1, 2019," NST-EST2019-alldata, 2019. Accessed: Dec. 06, 2020. [Online]. Available: <https://www.census.gov/data/tables/time-series/demo/popest/2010s-state-total.html>.
- [7] "WHO Coronavirus Disease (COVID-19) Dashboard." <https://covid19.who.int> (accessed Oct. 22, 2020).
- [8] "Microsoft Bing COVID-19 Tracker." http://bing.com/covid/local/maine_unitedstates?dynamicSharing=true&shp=Facebook&shw=900&shh=800&shtk=Y29yb25hdmlzdXMgdHJhY2tldiB1cGRhdGVz&shdk=dGVzdA%3D%3D&shth=OSH.Mmq%2BwuM5WWI/TcdViNGxBA&redire

- ct_uri=http%3A//veeraux%3A81/covid/local/unitedstates%3FdynamicSharing%3Dtrue&ref=Coronavirus&a1 (accessed Oct. 22, 2020).
- [9] "Track Testing Trends," *Johns Hopkins Coronavirus Resource Center*. <https://coronavirus.jhu.edu/testing/tracker> (accessed Nov. 17, 2020).
- [10] "COVID-19 Map," *Johns Hopkins Coronavirus Resource Center*. <https://coronavirus.jhu.edu/map.html> (accessed Oct. 22, 2020).
- [11] "The COVID Tracking Project Visualization Guide," *The COVID Tracking Project*. <https://covidtracking.com/about-data/visualization-guide> (accessed Oct. 22, 2020).
- [12] K. Field, "Mapping coronavirus, responsibly," *ArcGIS Blog*, Feb. 25, 2020. <https://www.esri.com/arcgis-blog/products/product/mapping/mapping-coronavirus-responsibly/> (accessed Dec. 03, 2020).
- [13] "A complete guide to coronavirus charts: Be informed, not terrified." <https://www.fastcompany.com/90477393/a-complete-guide-to-coronavirus-charts-be-informed-not-terrified> (accessed Oct. 22, 2020).
- [14] L. C. Rost, "17 (or so) responsible live visualizations about the coronavirus, for you to use," *Chartable*, Mar. 06, 2020. <https://blog.datawrapper.de/coronaviruscharts/index.html> (accessed Oct. 22, 2020).
- [15] S. Berinato, "Visualizations That Really Work," *Harvard Business Review*, no. June 2016, Jun. 01, 2016.
- [16] A. Makulec, "Ten Considerations Before you Create another Chart about COVID-19," *Medium*, Apr. 27, 2020. <https://medium.com/nightingale/ten-considerations-before-you-create-another-chart-about-covid-19-27d3bd691be8> (accessed Oct. 22, 2020).
- [17] M. D. Wilkinson *et al.*, "The FAIR Guiding Principles for scientific data management and stewardship," *Sci Data*, vol. 3, no. 1, p. 160018, Dec. 2016, doi: 10.1038/sdata.2016.18.
- [18] "EvanMPeck on Twitter," *Twitter*. <https://twitter.com/EvanMPeck/status/1235568532840120321> (accessed Oct. 22, 2020).

APPENDIX

Table A1. Distribution of work across team members

Task	Time Invested (hrs)		
	Cloris	Derek	Harry
Proposal	20	20	20
Analysis of dashboards	30	-	-
Analysis of design guidelines	9	4	-
Data preparation and analysis	-	-	20
Design brainstorming	4	4	4
Implementation	-	40	10
Project update write-up & meeting	15	15	15
Final presentation	8	10	10
Final report write-up	11	10	10